

Generating random thermal momenta

Denes Molnar

Physics Department, Purdue University, 525 Northwestern Avenue, West Lafayette, IN 47907, USA

(Dated: December 11, 2012)

Generation of random thermal particle momenta is a basic task in many problems, such as microscopic studies of equilibrium and transport properties of systems, or the conversion of a fluid to particles. In heavy-ion physics, the (in)efficiency of the algorithm matters particularly in hybrid *hydrodynamics + hadronic transport* calculations. With popular software packages, such as UrQMD 3.3p1 [1] or THERMINATOR [2], it can still take ten hours to generate particles for a single Pb+Pb “event” at the LHC from fluid dynamics output. Below I describe reasonably efficient basic algorithms using the MPC package (Molnar’s Program Collection) [3], which should help speed momentum generation up by at least one order of magnitude.

It is likely that this wheel has been reinvented many times instead of reuse, so there may very well exist older and/or better algorithms that I am not aware of (MPC has been around only since 2000). The main goal here is to encourage practitioners to use available efficient routines, and offer a few practical solutions.

A. Static thermal distributions

First consider a static thermal distribution

$$f_0(\mathbf{p}) \equiv \frac{dN}{d^3p} = \frac{g}{(2\pi)^3} \frac{1}{e^{(E(\mathbf{p})-\mu)/T} + a} \quad , \quad E(\mathbf{p}) = \sqrt{\mathbf{p}^2 + m^2} \quad (1)$$

where μ is the chemical potential, m is the particle mass, g is the degeneracy factor (internal degrees of freedom), and

$$a = \begin{cases} +1 & \text{for fermions} \\ -1 & \text{for bosons} \\ 0 & \text{for classical (Boltzmann) particles} \end{cases} \quad (2)$$

The distribution is isotropic, so in standard spherical coordinates we have *uniformly* distributed

$$\varphi \in [0, 2\pi) \quad , \quad \cos \vartheta \in [-1, 1] \quad (3)$$

and the only challenge is to generate the scaled magnitude $x = |\mathbf{p}|/T$

$$g(x) \equiv \frac{dN}{dx} \propto \frac{x^2}{e^{\sqrt{x^2+z^2}} + \alpha} \quad \left(z \equiv \frac{m}{T} \quad , \quad \alpha \equiv a e^{\mu/T} \right) \quad (4)$$

A straightforward technique to employ is an automated version of the *rejection method* (see Fig. 1). The rejection method[4] is based on a suitable comparison function F that bounds g from above for all x . One then generates uniformly (x, y) pairs in the two-dimensional area under the graph of F , and keeps the x coordinate of those points that are below the graph of g . It is easy to construct a *staircase* comparison function, automatically, via tabulating $g(x)$ at several points, and multiplying by a modest “safety” factor to leave room for any local maxima within the intervals [5]. Uniform sampling under the graph of a staircase function is straightforward[7][5].

Alternatively, we can tabulate $g(x)$, construct an *interpolation* of it, and generate the interpolated distribution[8]. Computationally this is very similar to staircase function sampling in automated rejection. To precisely represent $g(x)$, in general many points are needed but we then have 100% efficiency (there is no rejection check).

B. Boosted thermal distributions

For thermal distribution with collective flow

$$f_v(\mathbf{p}) \equiv \frac{dN}{d^3p} = \frac{g}{(2\pi)^3} \frac{1}{e^{(p^\alpha u_\alpha - \mu)/T} + a} \quad , \quad p^0 = \sqrt{\mathbf{p}^2 + m^2} \quad , \quad u^\mu = \gamma(1, \mathbf{v}_F) \quad (5)$$

where \mathbf{v}_F is the three-velocity of the fluid. This is a special case ($n^\mu = (1, \mathbf{0})$) of Cooper-Frye freezeout discussed in the next Section.




FIG. 1: Illustration of rejection method with staircase comparison function.

C. Cooper-Frye freezeout

In Cooper-Frye freezeout[6] a fluid is converted to particles across a three-dimensional spacetime hypersurface. A surface element $d\sigma^\mu(x)$ at spacetime point x contributes to particle number

$$dN = p^\alpha d\sigma_\alpha(x) \frac{d^3p}{p^0} f(x, \mathbf{p}) \quad (d\sigma_\alpha \equiv n_\alpha d^3V, \quad n^\alpha = (n^0, \mathbf{n}), \quad n^2 = \pm 1) \quad (6)$$

where f is the phase space density of the particles in the fluid. For an ideal fluid f is thermal, so the momentum distribution of particles emitted from the surface element is

$$f_{CF}(\mathbf{p}) = \frac{dN}{d^3p} \propto \frac{p_\alpha n^\alpha}{p^0} f_v(\mathbf{p}) \quad (7)$$

with f_v from (5). Because dN is a Lorentz scalar, in the local rest frame of the fluid element ($\tilde{\mathbf{v}}_F = \mathbf{0}$) we have a static thermal distribution with a momentum-dependent prefactor

$$\frac{dN}{d^3p} \propto \frac{\tilde{p}_\alpha \tilde{n}^\alpha}{\tilde{p}^0} f_0(\tilde{\mathbf{p}}) = (\tilde{n}^0 - \tilde{\mathbf{n}}\tilde{\mathbf{v}}) f_0(\tilde{\mathbf{p}}), \quad \tilde{\mathbf{v}} \equiv \frac{\tilde{\mathbf{p}}}{\tilde{p}^0} \quad (8)$$

After generating a thermal momentum $\tilde{\mathbf{p}}$, the prefactor can be accounted for in an additional rejection step. If $\tilde{\mathbf{p}}$ fails the rejection test, we generate a new one, otherwise we boost it back to the original frame to obtain \mathbf{p} .

For *timelike* normal with $\tilde{n}^0 > 0$, the prefactor is always nonnegative, and the additional rejection step is at least 50% efficient because we can use as comparison function $\tilde{n}^0 + |\tilde{\mathbf{n}}| < 2\tilde{n}^0$, while on average $\langle \tilde{n}^0 - \tilde{\mathbf{n}}\tilde{\mathbf{v}} \rangle = \tilde{n}^0$. Negative $\tilde{n}^0 < 0$ is unphysical (dN is never positive then). So we need no more than *two tries*, on average.

For *spacelike* normal ($n^2 = -1$), there are always momenta for which $(\tilde{n}^0 - \tilde{\mathbf{n}}\tilde{\mathbf{v}})$ is positive, and also momenta for which it is negative, so one needs to follow some prescription to handle unphysical contributions. A customary choice is to ignore negative values and implement a *cut* prefactor $(\tilde{n}^0 - \tilde{\mathbf{n}}\tilde{\mathbf{v}})\Theta(\tilde{n}^0 - \tilde{\mathbf{n}}\tilde{\mathbf{v}})$ instead, even if with thermal f_0 conservation laws are violated[9]. For $\tilde{n}^0 \geq 0$, a comparison function $\tilde{n}^0 + |\tilde{\mathbf{n}}|$ then provides[10] an efficiency of at least $\langle |\tilde{\mathbf{v}}| \rangle / 4$, which for $m/T \gg 1$ is poor but still workable in heavy-ion physics applications because spacelike surface elements are rare, and heavy particles have low abundances. For $\tilde{n}^0 < 0$ the total particle yield from the hypersurface element (without the cut) is negative, and for that reason such surface elements are most often omitted. Momentum generation for $\tilde{n}^0 < 0$, if desired, requires additional care[11].

For better efficiency, especially for spacelike normals and heavy particles, we can generate directly the isotropic

$$(\tilde{n}^0 + |\tilde{\mathbf{n}}||\tilde{\mathbf{v}}|) \Theta(\tilde{n}^0 + |\tilde{\mathbf{n}}||\tilde{\mathbf{v}}|) f_0(\tilde{\mathbf{p}}), \quad (9)$$

and realize the correct prefactor $(\tilde{n}^0 - \tilde{\mathbf{n}}\tilde{\mathbf{v}})\Theta(\tilde{n}^0 - \tilde{\mathbf{n}}\tilde{\mathbf{v}})$ via rejection using the more efficient[12] comparison function $\tilde{n}^0 + |\tilde{\mathbf{n}}||\tilde{\mathbf{v}}|$. This way, for *timelike* normals we need no more than *two tries*, on average, while for *spacelike* normals with $\tilde{n}^0 \geq 0$, no more than *four tries*. The price is that a new generator has to be constructed for each hypersurface element, whereas the static thermal generator (f_0) can be reused when $T = \text{const}$ across the entire hypersurface (often the case in practice).

D. Massless Boltzmann distribution

Finally, I show that for *massless* Boltzmann particles static thermal momenta

$$\frac{dN}{dx} \propto x^2 e^{-x} \quad , \quad x \equiv \frac{|\mathbf{p}|}{T} \quad , \quad (10)$$

can be generated through the one-liner[13]

$$x = -\ln[(1 - \xi_1)(1 - \xi_2)(1 - \xi_3)] \quad (11)$$

where $\xi_{1,2,3}$ are three independent uniform deviates on $[0, 1)$. This follows because with $x = r^2$, the distribution corresponds to a Gaussian in six-dimensional space $\{y_1, y_2, \dots, y_6\}$

$$\frac{dN}{r^5 dr} \propto e^{-r^2} \quad , \quad r = \sqrt{y_1^2 + y_2^2 + \dots + y_6^2} \quad (12)$$

which can be generated as three independent *pairs* of Gaussian deviates, each pair through the standard 2D polar coordinate technique $dx dy e^{-x^2} e^{-y^2} \propto d\varphi d(\rho^2) e^{-\rho^2}$. Here φ is uniform on $[0, 2\pi)$ and ρ^2 is exponential on $[0, \infty)$. We only need ρ^2 for each pair because $r^2 = \rho_{12}^2 + \rho_{34}^2 + \rho_{56}^2$, and exponential deviates can be generated through inverting

$$I(\rho^2) = \int_0^{\rho^2} dw e^{-w} = 1 - e^{-\rho^2} \quad \Rightarrow \quad \rho^2 = -\ln(1 - I) \quad (13)$$

where I is a uniform deviate on $[0, 1)$.

E. Conclusion

In a paradigm where random numbers are computationally very expensive, a useful efficiency benchmark for three-momentum generators is

$$\epsilon = \frac{3}{\langle \text{uniform deviates per momentum} \rangle} \quad (14)$$

(i.e., with one random number generator call per momentum component, $\epsilon = 1$). For the algorithms above - static thermal (Sec. A), boosted thermal (Sec. B), and Cooper-Frye (Sec. C) - $\epsilon = 1$, $\epsilon \geq 3/8$, and $\epsilon \geq 3/16$, respectively, if we use the interpolation method[14].

In reality there is of course overhead due to function tabulation, memory lookups, and binary search, which limit the number of intervals one can utilize. But in practice, overall efficiency is still a reasonable ~ 0.1 (even with the automated rejection method), and often much better[15].

Acknowledgments. - Discussions with Giorgio Torrieri, Pasi Huovinen, and Sangwook Ryu are greatly acknowledged. The routines described here were developed under DOE grants DE-FG-02-93ER-40764 (Columbia), DE-FG02-01ER41190 (Ohio State), DE-AC02-98CH10886 (RIKEN BNL), and DE-PS02-09ER41665 (Purdue).

- [1] H. Petersen, J. Steinheimer, G. Burau, M. Bleicher and H. Stöcker, Phys. Rev. C **78**, 044901 (2008) [arXiv:0806.1695 [nucl-th]].
- [2] A. Kisiel, T. Taluc, W. Broniowski and W. Florkowski, Comput. Phys. Commun. **174**, 669 (2006) [nucl-th/0504047].
- [3] The MPC package can be downloaded from the Open Standard Codes And Routines (OSCAR) repository at <http://karman.physics.purdue.edu/OSCAR>
- [4] See, e.g., Chapter 7 of W. H. Press *et al*, *Numerical Recipes in C*, 2nd Ed. (Cambridge University Press, 2002)
- [5] See, e.g., the `AutoRejection1D` class in MPC.
- [6] F. Cooper and G. Frye, Phys. Rev. D **10** (1974) 186.
- [7] E.g., generate a uniform deviate I , and find the interval $[x_i, x_{i+1}]$ that has the point x satisfying $I = P(x) \equiv \int_0^x dx' F(x') / \int_0^\infty dx' F(x')$. With precalculated $P(x_i)$ this boils down to a binary search. Then pick y uniformly in $[0, F(x)]$.

- [8] See `TabulatedRndDist1D` class in MPC.
- [9] Unless we also alter the hydrodynamic variables that define the distribution. See, e.g., C. Anderlik *et al*, Phys. Rev. C **59**, 388 (1999) [nucl-th/9808024].
- [10] If $\tilde{n}^0 \geq |\tilde{\mathbf{n}}||\tilde{\mathbf{v}}|$, the Θ -function is unity and the efficiency has the same expression as for a timelike normal $\tilde{n}^0/(\tilde{n}^0 + |\tilde{\mathbf{n}}|) \geq |\tilde{\mathbf{v}}|/2$ because $|\mathbf{n}| > \tilde{n}^0$. While if $0 \leq \tilde{n}^0 < |\tilde{\mathbf{n}}||\tilde{\mathbf{v}}|$, on average $\langle(\tilde{n}^0 - \tilde{\mathbf{n}}\tilde{\mathbf{v}})\Theta(\tilde{n}^0 - \tilde{\mathbf{n}}\tilde{\mathbf{v}})\rangle = (\tilde{n}^0 + |\tilde{\mathbf{n}}||\tilde{\mathbf{v}}|)^2/(4|\tilde{\mathbf{n}}||\tilde{\mathbf{v}}|) \geq (\tilde{n}^0 + |\tilde{\mathbf{n}}||\tilde{\mathbf{v}}|)(\tilde{n}^0 + |\tilde{\mathbf{n}}|)/(4|\tilde{\mathbf{n}}|)$, so the efficiency is at least $(\tilde{n}^0 + |\tilde{\mathbf{n}}||\tilde{\mathbf{v}}|)/(4|\tilde{\mathbf{n}}|) \geq |\tilde{\mathbf{v}}|/4$.
- [11] When $|\tilde{\mathbf{n}}|$ is close to $|\tilde{n}^0|$, only high-velocity particles moving in nearly opposite direction to $\tilde{\mathbf{n}}$ can contribute because otherwise the prefactor is negative, and a naive rejection algorithm can turn into a near-infinite loop. For reasonable efficiency one must generate *directly* only the allowed tail and limited solid angle of the static thermal distribution f_0 or of distribution (9). The constraints are i) $|\tilde{\mathbf{v}}| \geq |\tilde{n}^0|/|\tilde{\mathbf{n}}|$, i.e., $|\tilde{\mathbf{p}}| \geq m|\tilde{n}^0|$; and ii) the angle between $\tilde{\mathbf{n}}$ and $\tilde{\mathbf{p}}$ satisfies $\cos \tilde{\vartheta} \leq -|\tilde{n}^0|/|\tilde{\mathbf{n}}|$.
- [12] In particular, for spacelike normals with $\tilde{n}^0 \geq 0$, the efficiency is $\langle(\tilde{n}^0 - \tilde{\mathbf{n}}\tilde{\mathbf{v}})\Theta(\tilde{n}^0 - \tilde{\mathbf{n}}\tilde{\mathbf{v}})\rangle/(\tilde{n}^0 + |\tilde{\mathbf{n}}||\tilde{\mathbf{v}}|) = (\tilde{n}^0 + |\tilde{\mathbf{n}}||\tilde{\mathbf{v}}|)/(4|\tilde{\mathbf{n}}||\tilde{\mathbf{v}}|) \geq 1/4$.
- [13] Routine `boltzmandev0` in MPC.
- [14] With automated rejection the formal efficiency drops to $\epsilon \gtrsim 3/4$, $\epsilon \gtrsim 3/10$, $\epsilon \gtrsim 3/20$ because with a nearly perfect comparison function we still need two random deviates per call.
- [15] The main inefficiency in UrQMD and THERMINATOR comes from the use of a *constant* comparison function, directly for the three-dimensional $f(\mathbf{p})$.